



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

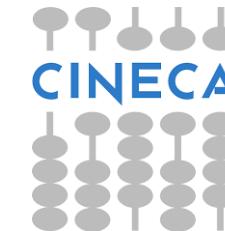
: Combine Out-of-band monitoring with AI and big data for datacenter automation



EU H2020 FETHPC
project ANTAREX
(g.a. 671623)

ETH zürich

Andrea Bartolini, Prof
University of Bologna – DEI, Italy



EU FP7 ERC Project
MULTITHERMAN
(g.a.291125)

E4
COMPUTER
ENGINEERING



The University of Bologna

THE BIRTH OF THE UNIVERSITY

The Studium in Bologna is the first home of free teaching, independent from ecclesiastic schools. Irnerio's law school marks the birth of Western universities.



FREEDOM OF RESEARCH

Federico I Barberino grants the University as the first of masters and professors the right to undertake to practice their profession travelling for the study. For the first time absolute freedom of research is ratified.

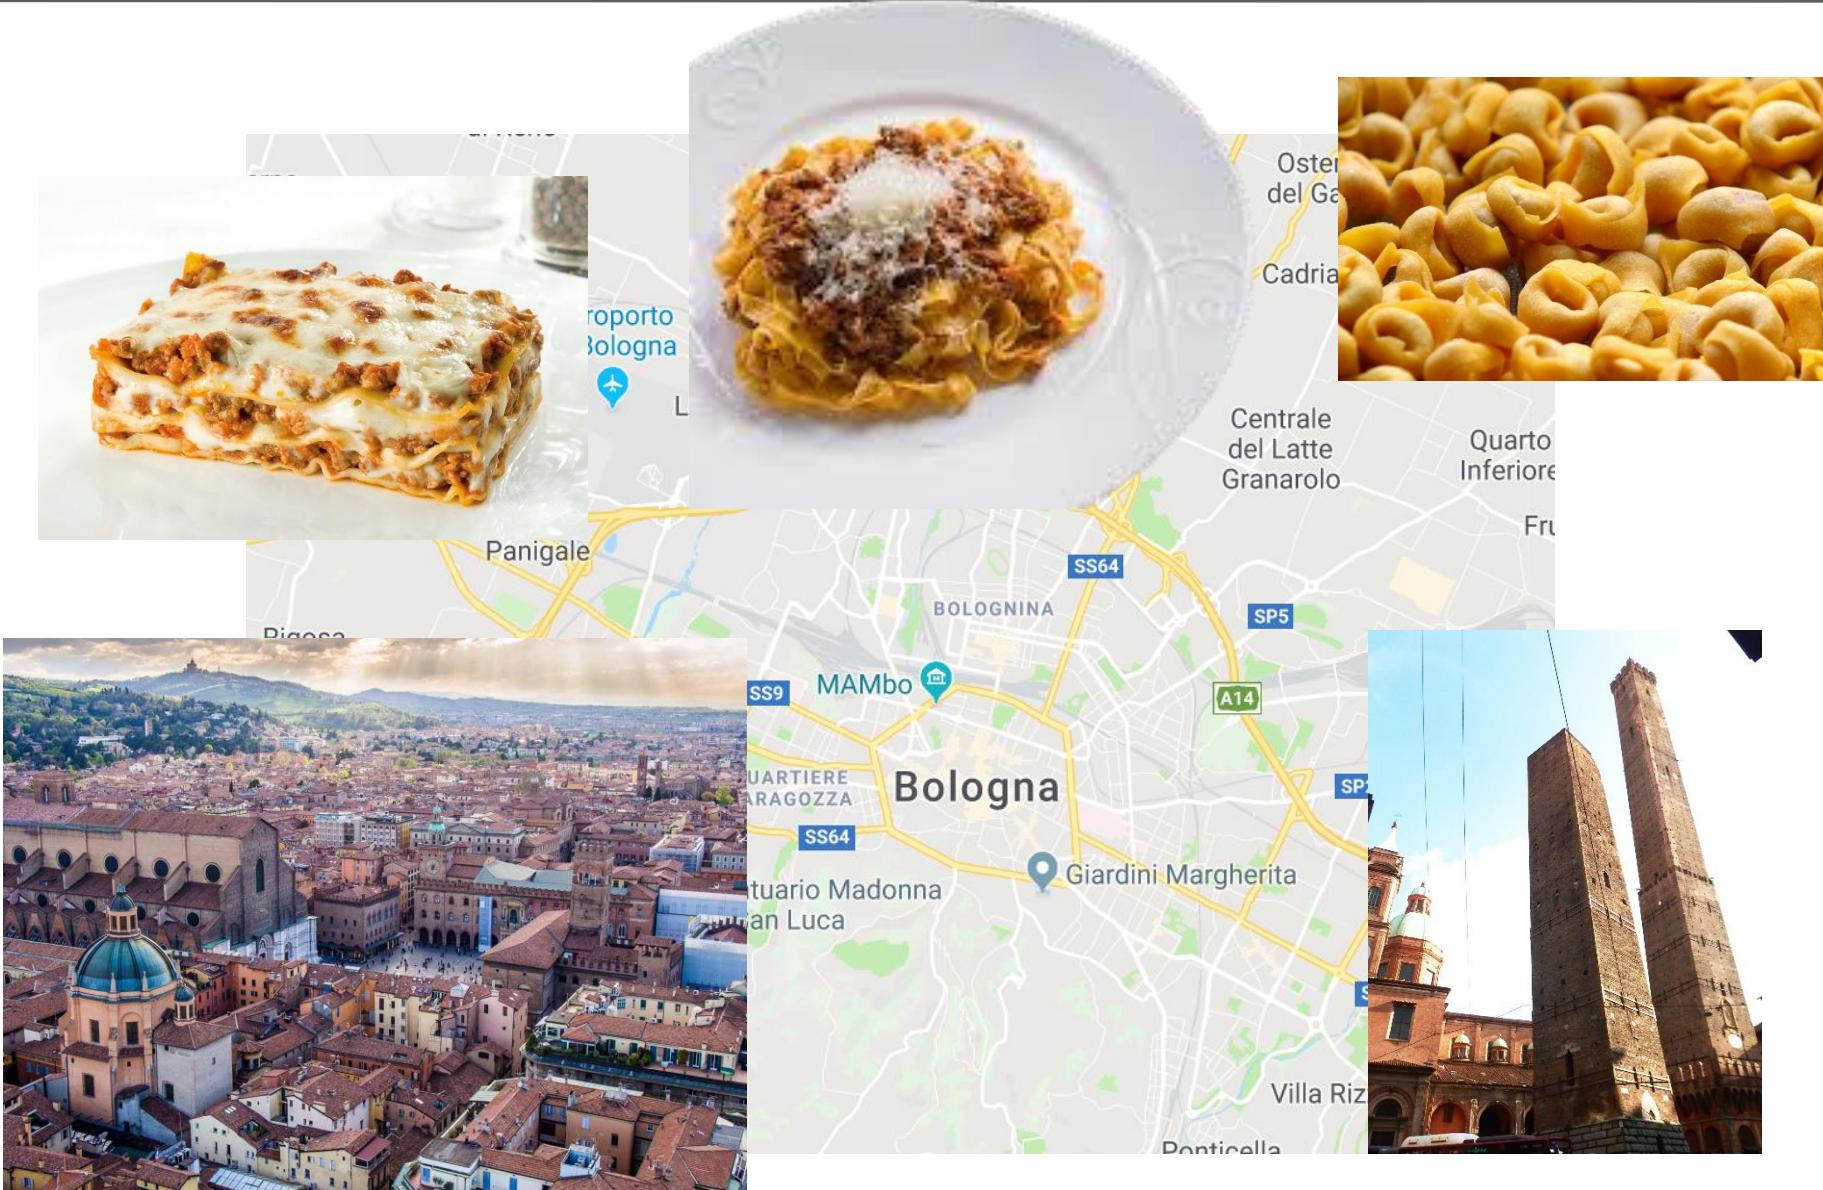
Alma Mater Studiorum
Università di Bologna
is a multi-campus
university based in Bologna,
Cesena, Forlì, Ravenna, and
Rimini.

University Statute, Constituent
Principles, Art. 1 para. 2





Bologna the city of ...





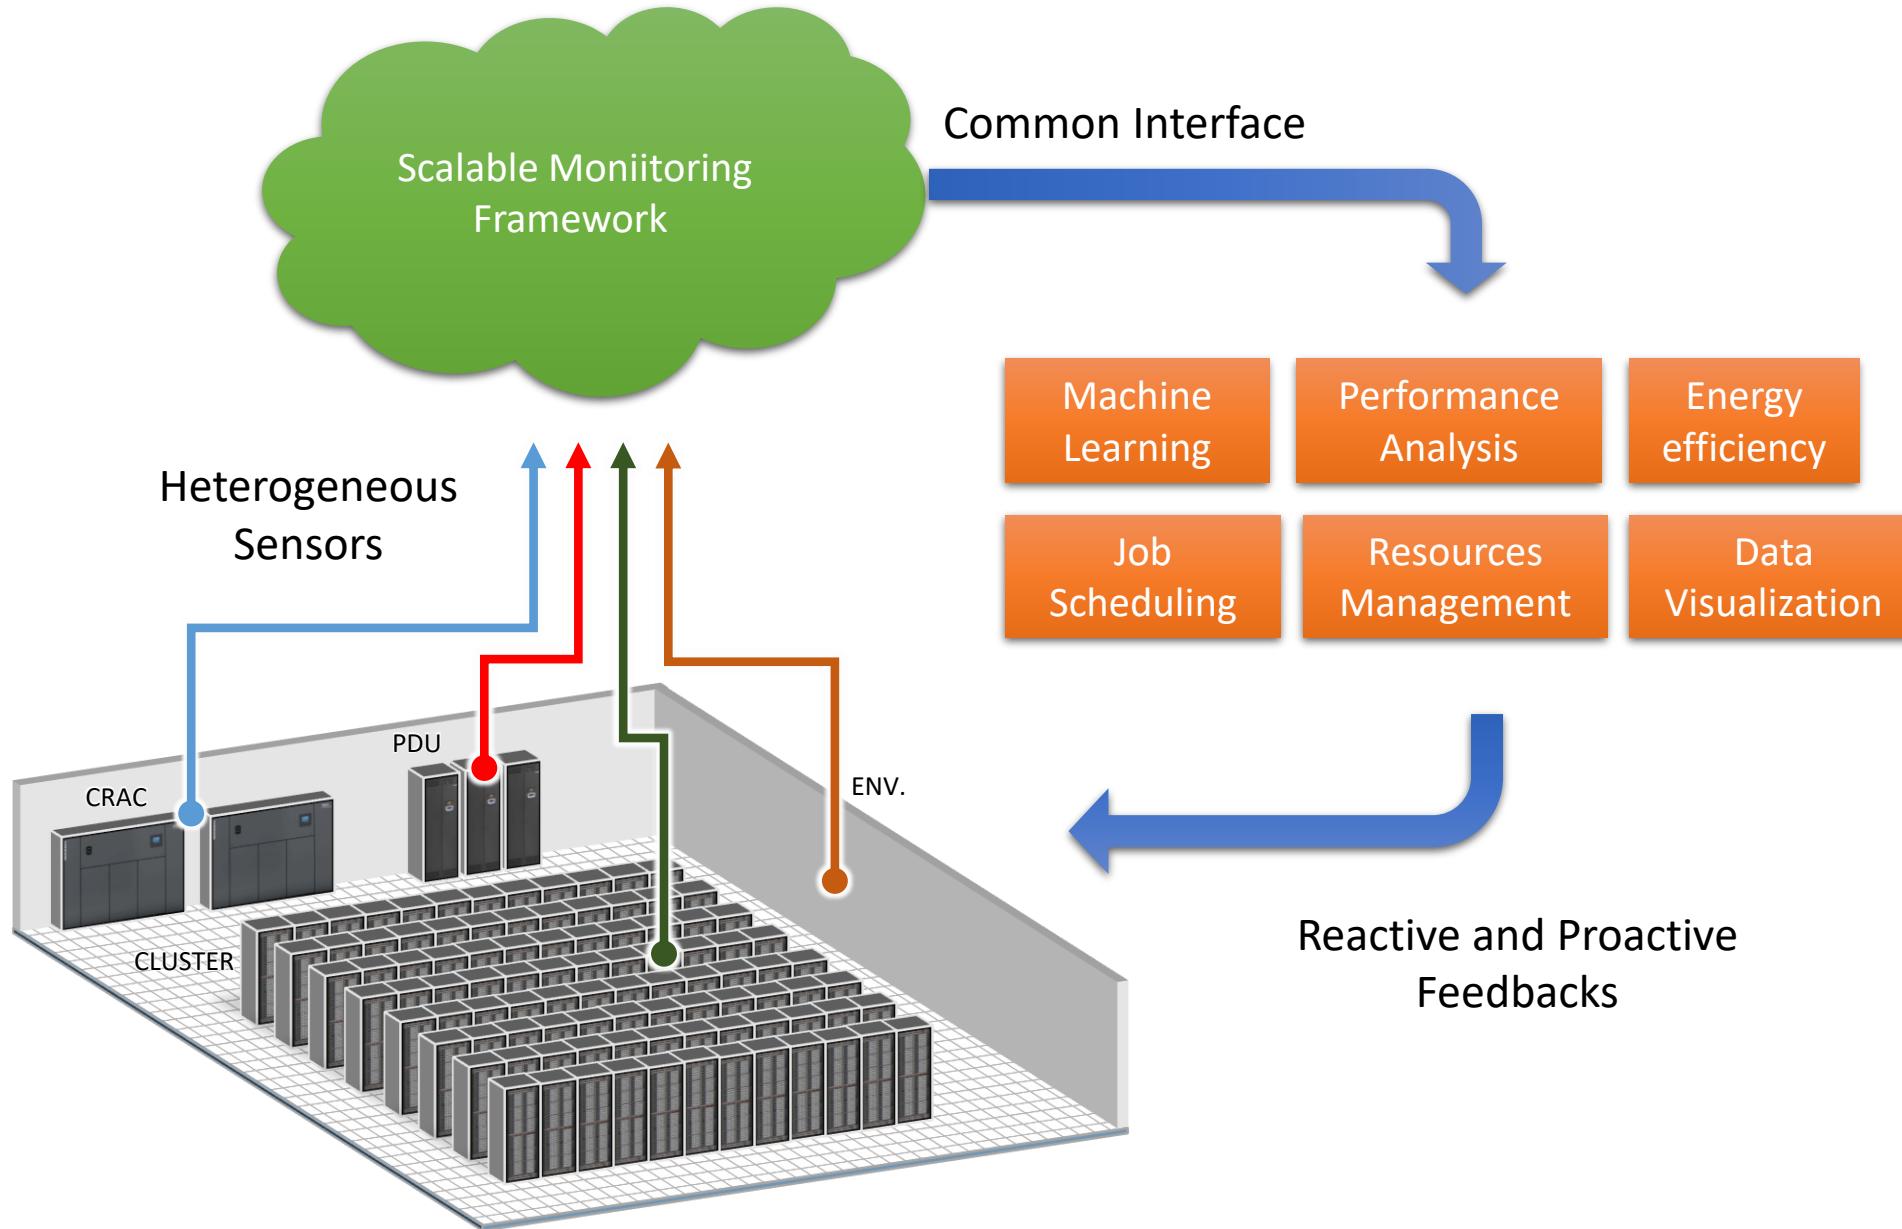
Bologna the city of ...



...Datacentres.



A New Trend: Datacentre Automation

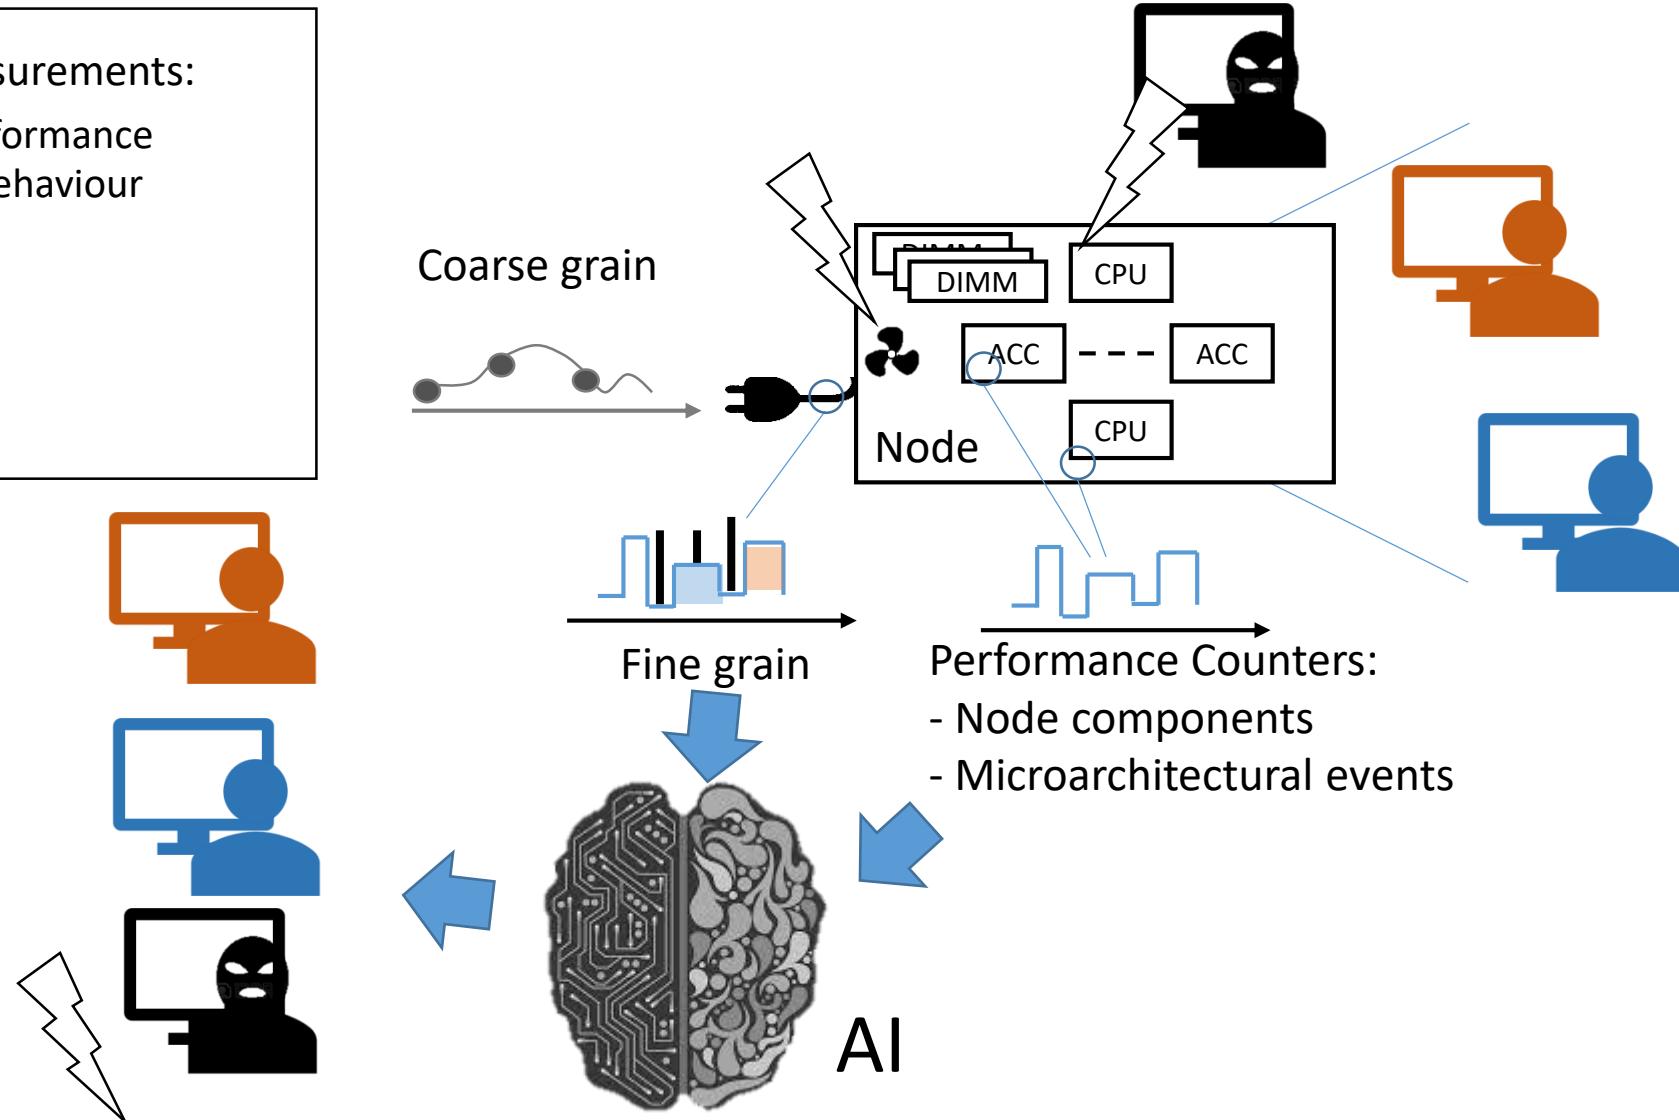




Usage Scenarios

Fine Grain Power and Performance Measurements:

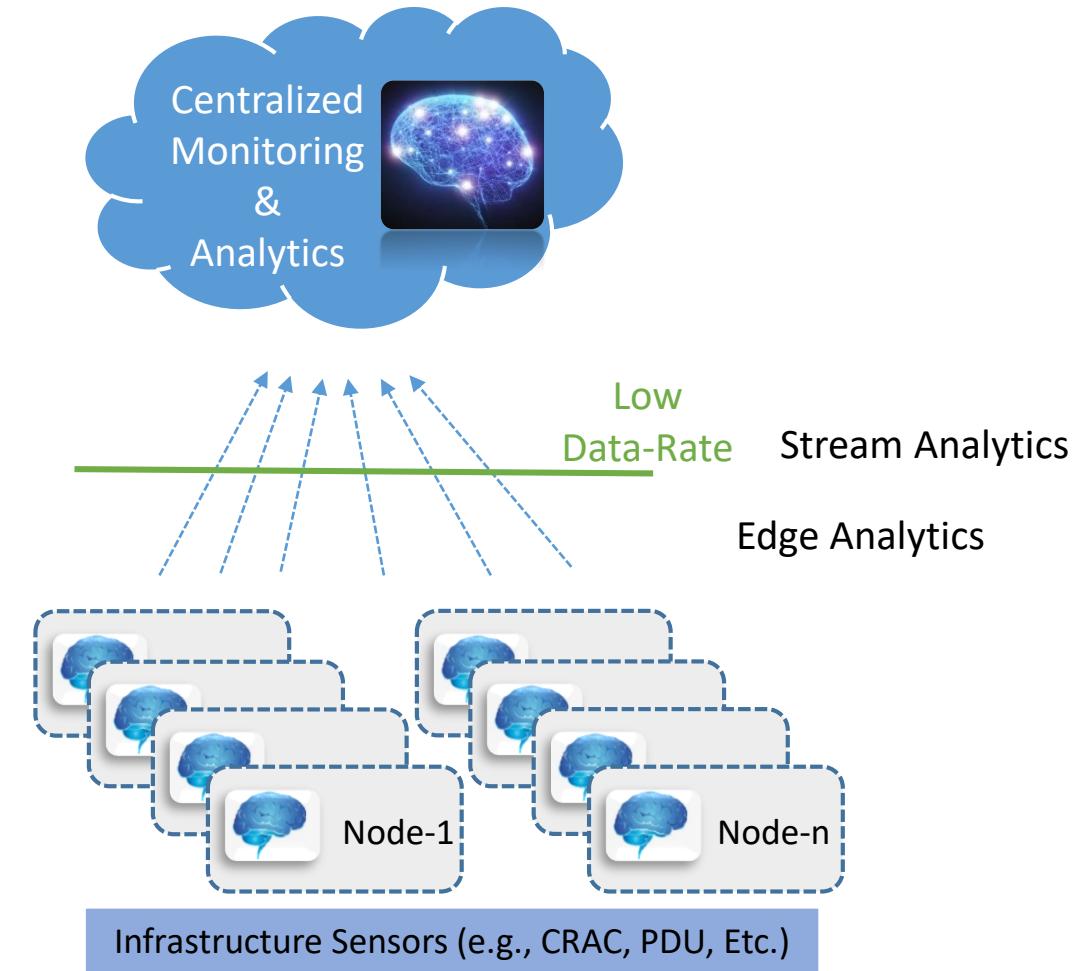
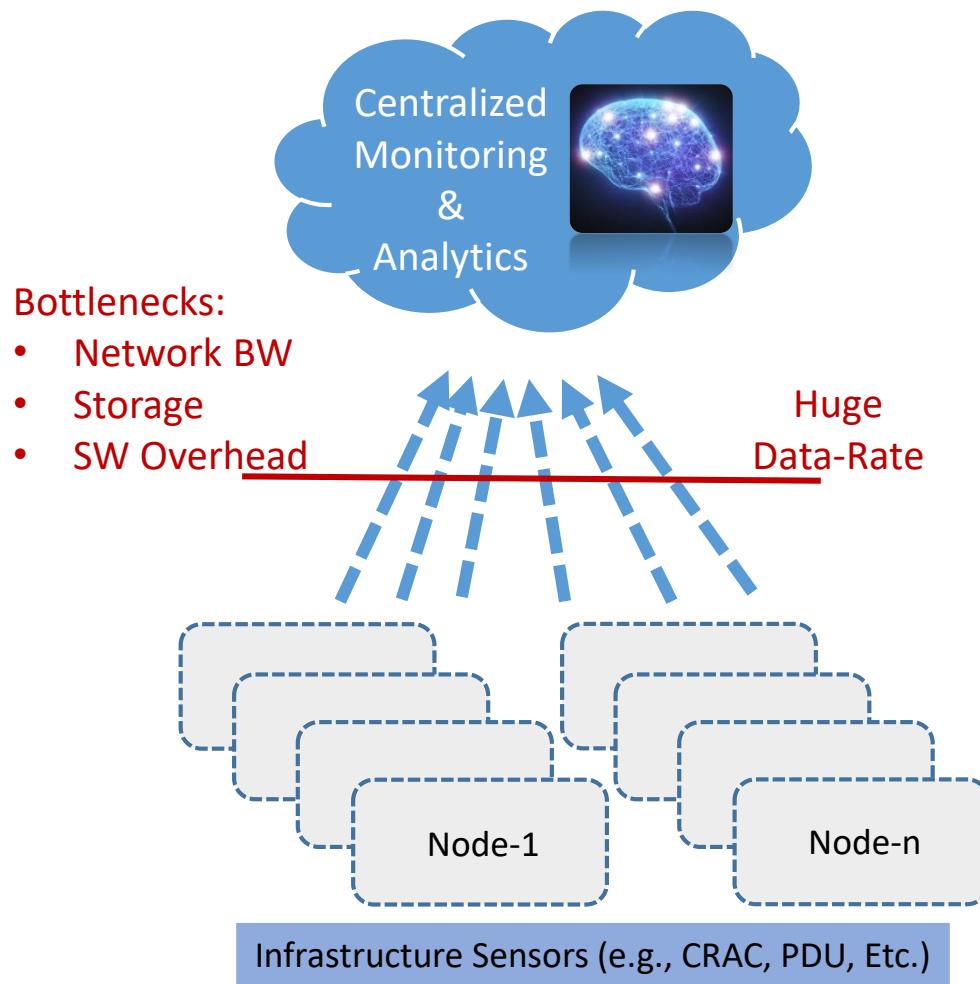
- Verify and classify node performance
 - In spec / out of spec behaviour
 - Miss configuration
 - Aging and wear out
- Detect security hazards
- Predictive maintenance





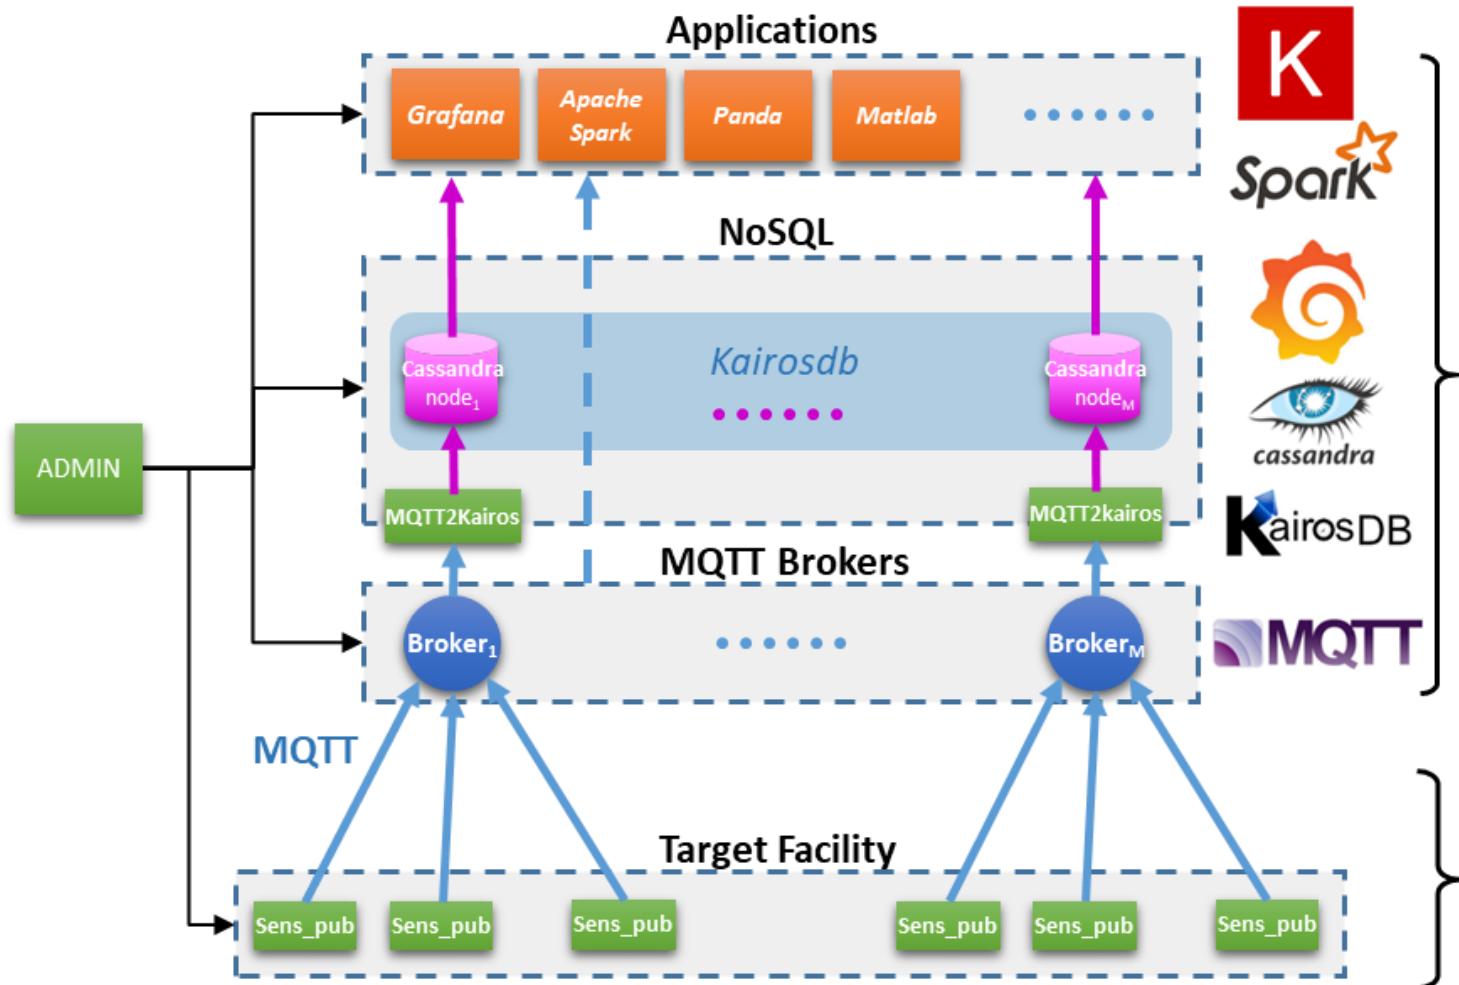
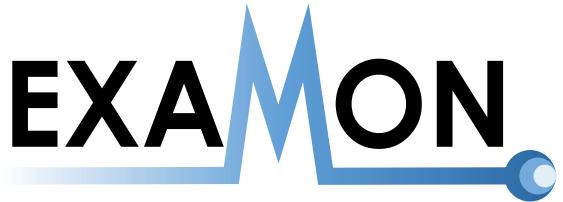
Datacenter Automation Design and Bottlenecks

[DAAC18] Libri et al. 2018. DiG: Enabling Out-of-Band Scalable High-Resolution Monitoring for Data-Center Analytics, Automation and Control





Scalable Data Collection, Analytics



Front-end

- MQTT Brokers
- Data Visualization
- NoSQL Storage
- Big Data Analytics

Back-end

- MQTT-enabled sensor collectors

<https://github.com/EEESlab/examon>

[DATE17] F. Beneventi et al., “Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools”
[BigDAW18] A. Bartolini et al., “The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support”

DiG = High Frequency Monitoring on D.A.V.I.D.E.



wistron®
CINECA



OCP form factor compute node
based on IBM Minsky

4x NVIDIA Tesla
P100 HSMX2

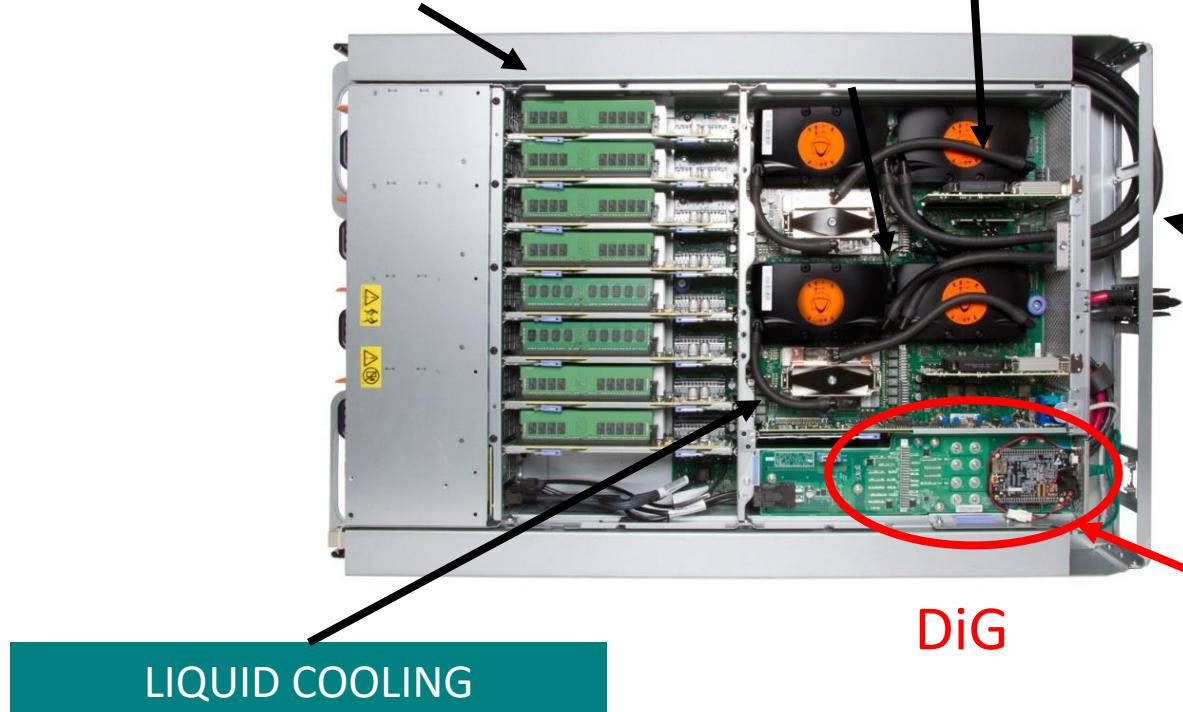
2 x IBM POWER8
with NVLink

2xIB EDR

BusBar

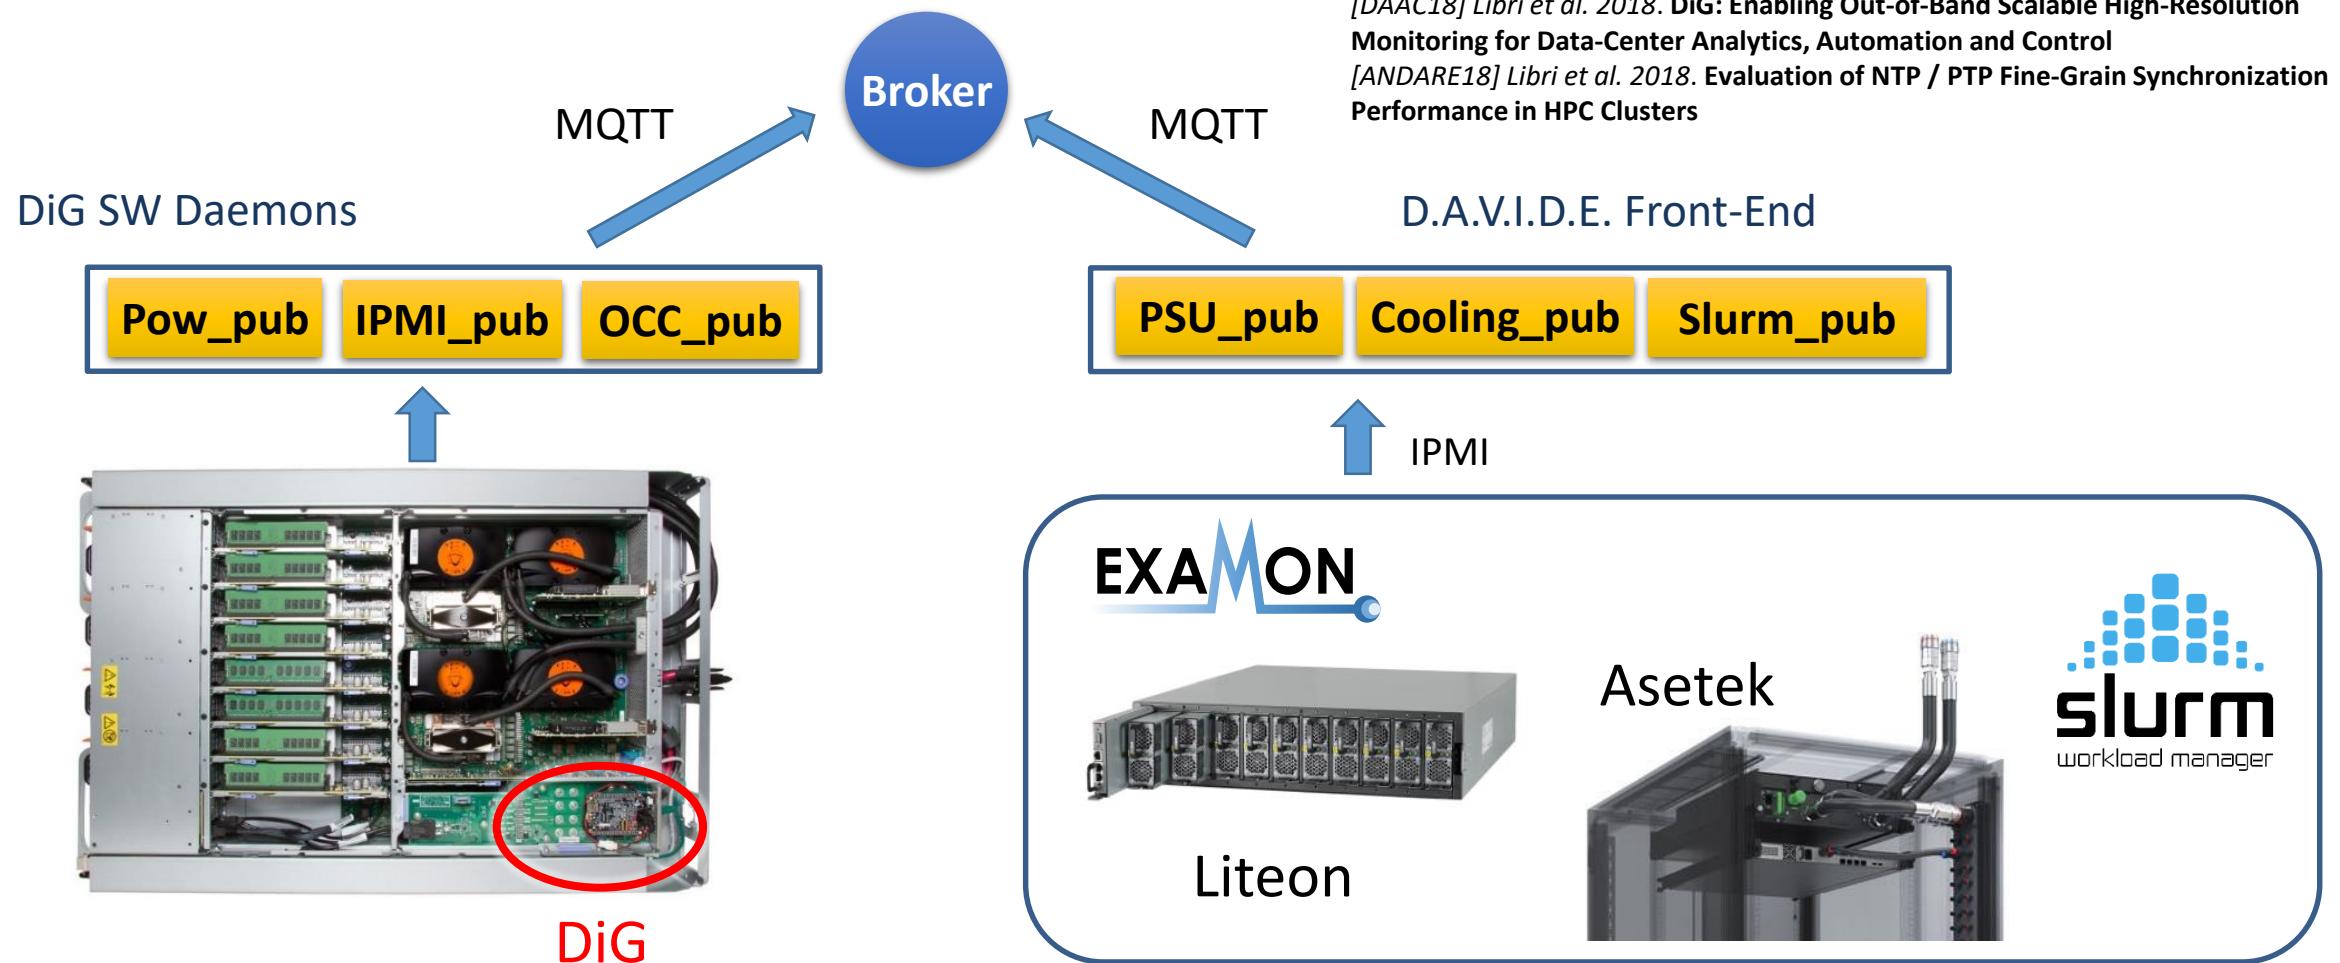
DiG

ETH Zurich / Univ. of Bologna
SoA out-of-band
High Resolution Power
and Performance Monitoring



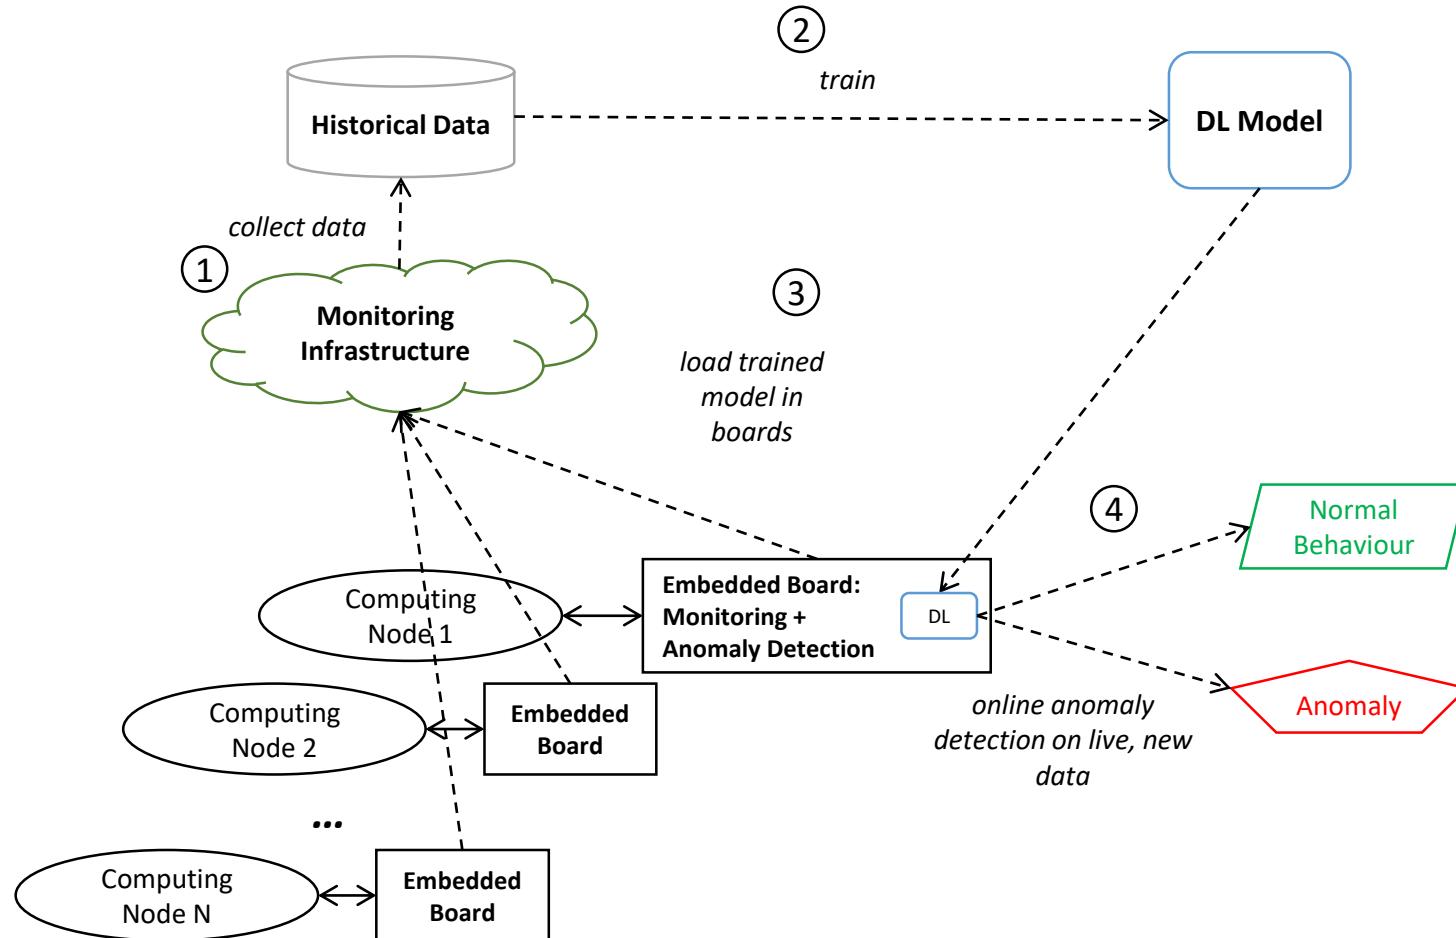
[DAAC18] Libri et al. 2018. DiG: Enabling Out-of-Band Scalable High-Resolution Monitoring for Data-Center Analytics, Automation and Control [BigDAW18] A. Bartolini et al, "The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support"

DiG = High Frequency Monitoring on D.A.V.I.D.E.





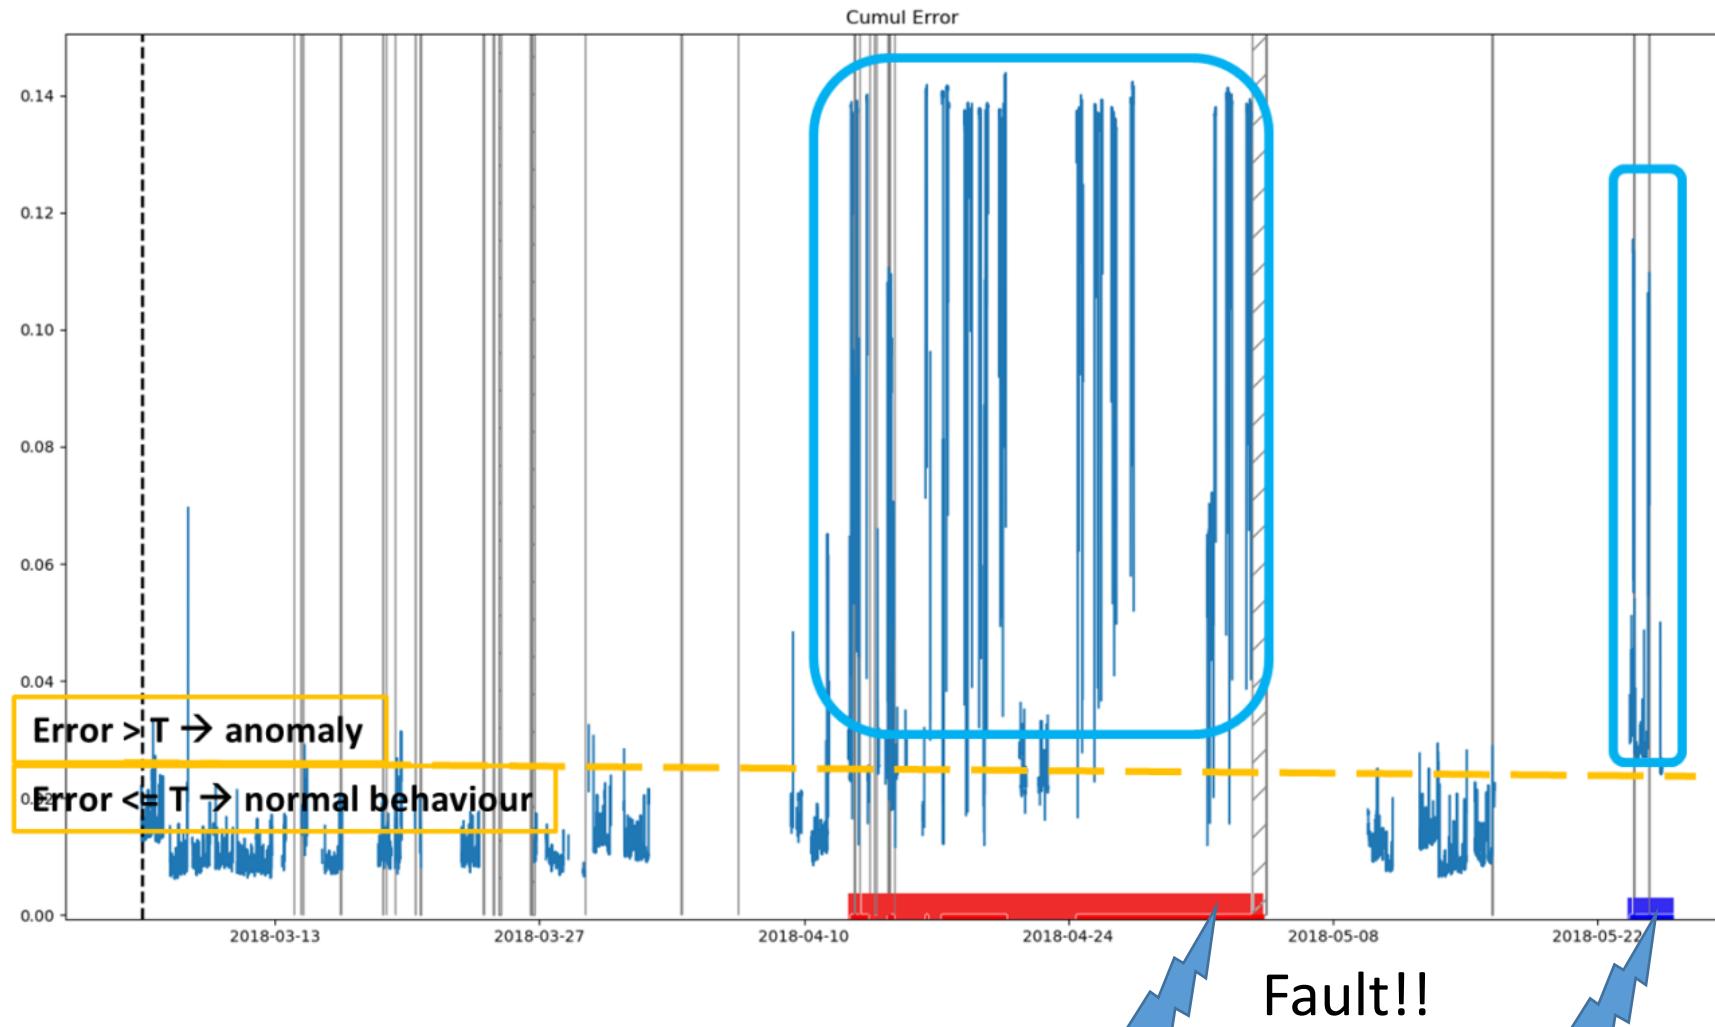
AI+Big Data on D.A.V.I.D.E.: Anomaly detection





AI+Big Data on D.A.V.I.D.E.: Anomaly detection

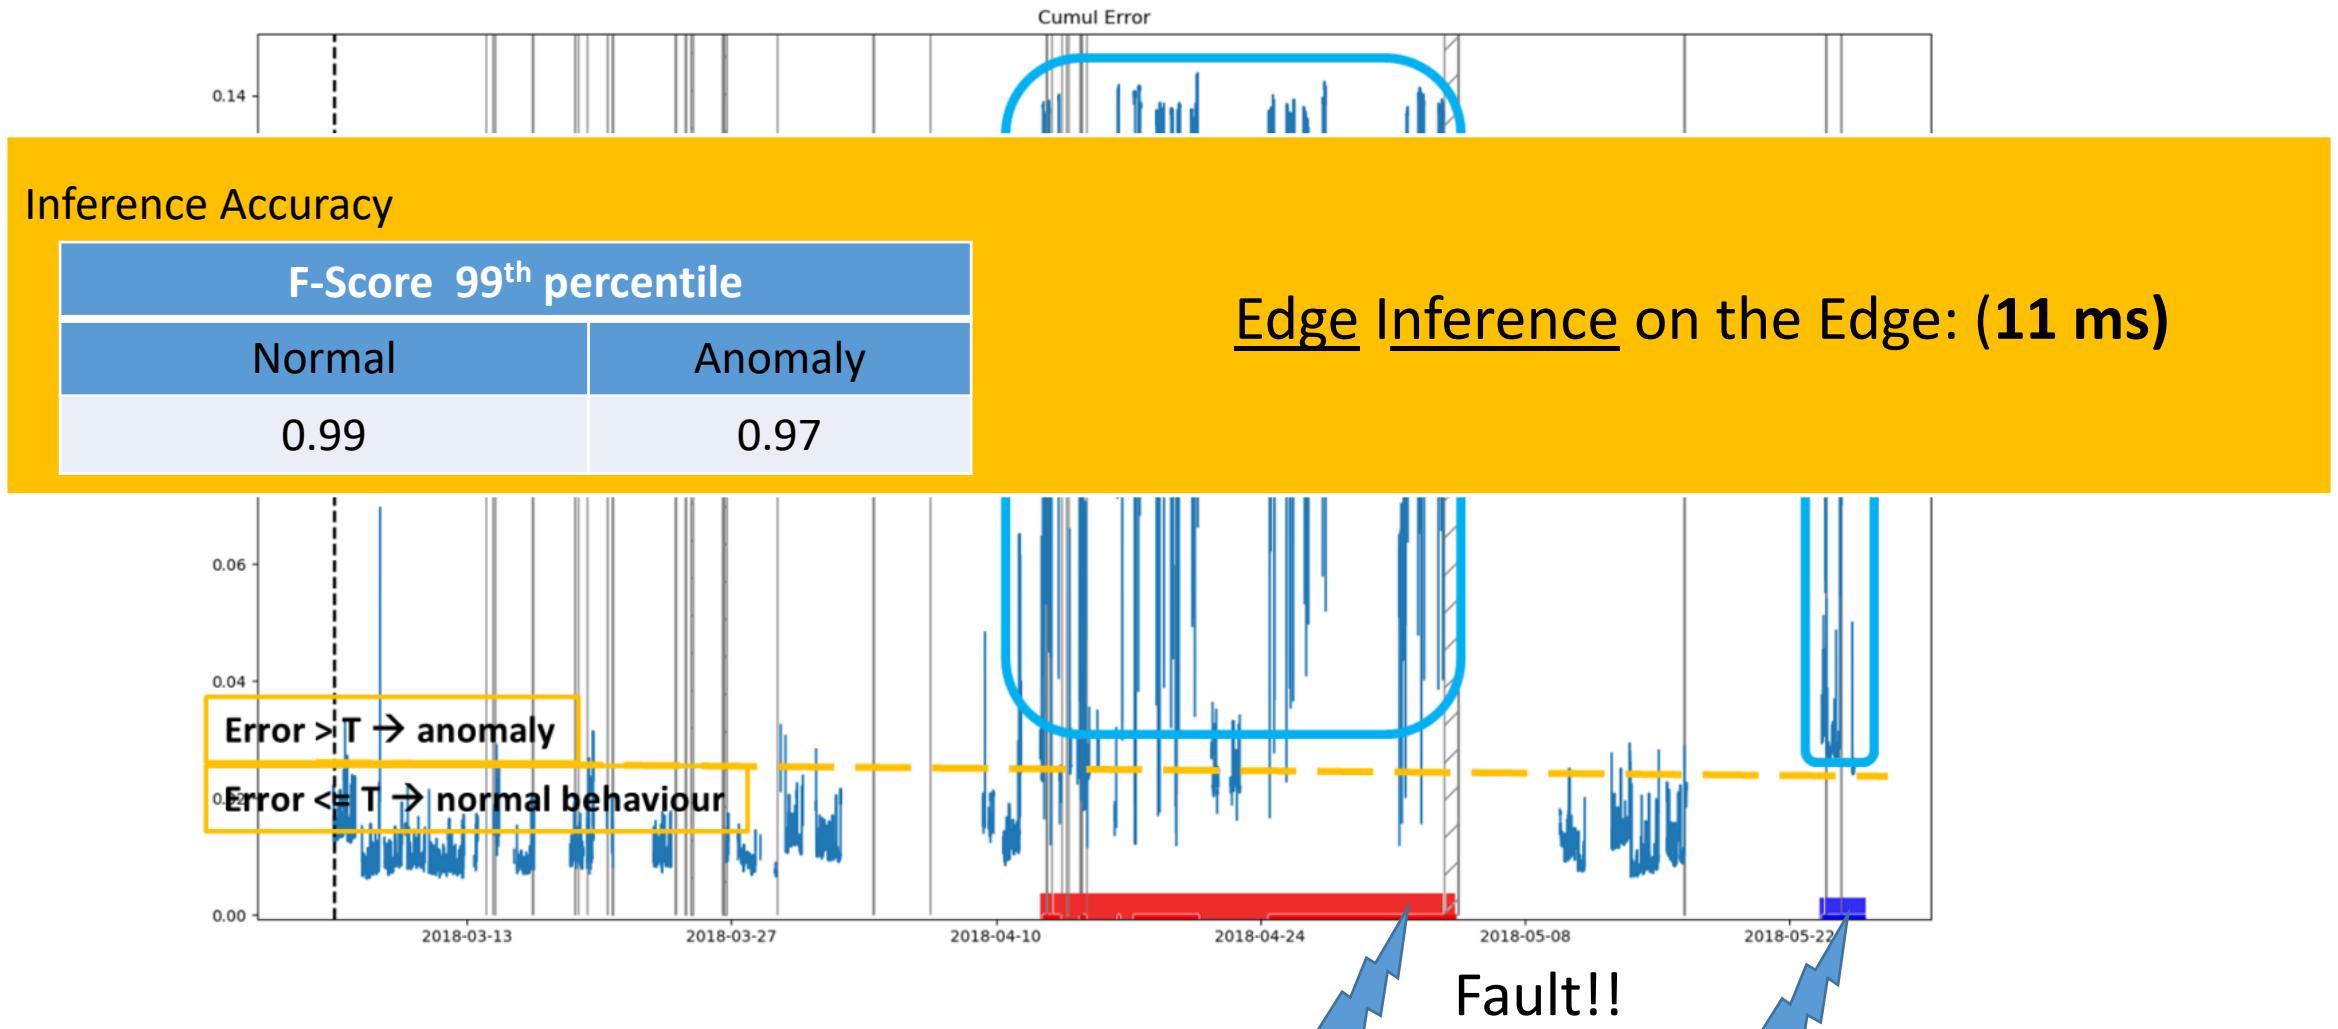
[IAAI19] Borghesi et al, <https://arxiv.org/abs/1811.05269>





AI+Big Data on D.A.V.I.D.E.: Anomaly detection

[IAAI19] Borghesi et al, <https://arxiv.org/abs/1811.05269>





Conclusion & Future Works

- Holistic and Fine Grain Monitoring feasible w. Open Source and Scalable of the shelf tools
- Challenge in the effective usage and knowledge extraction from the monitored data
- AI, ML and Predictive Control can leverage the infrastructure toward datacentre automation.
- Future Works:
 - Scale anomaly detection at the full datacentre level
 - HW design support for Datacenter Automation



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA



EU H2020 FETHPC
project ANTAREX
(g.a. 671623)

ETH zürich

ACKNOWLEDGE

The UNIBO/ETHZ Datacenter Automation TEAM

- Luca Benini, Michela Milano, Andrea Tilli, Roberto Diversi, Cristian Conficoni, Andrea Borghesi, Daniele Cesarini Antonio Libri, Francesco Beneventi, Federico Pittino

CINECA:

- Sanzio Bassini, Carlo Cavazzoni, Elda Rossi, Daniela Galletti, Alessandro De Federico, Isabella Baccarelli, Marco Sbrighi, Giuseppa Muscianisi

E4:

- Cosimo Gianfreda, Fabrizio Magugliani, Daniele Gregori, Simone Tinti



Multitherman

EU FP7 ERC Project
MULTITHERMAN
(g.a.291125)

E4
COMPUTER
ENGINEERING



References

- F. Beneventi et al. 2017, “**Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools**”, DATE 17
- A. Bartolini et al 2018, “**The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support**”, BIGDAW 2018
- Libri et al. 2018. **Evaluation of NTP / PTP Fine-Grain Synchronization Performance in HPC Clusters**, ANDARE'18
- Libri et al. 2018. **DiG: Enabling Out-of-Band Scalable High-Resolution Monitoring for Data-Center Analytics, Automation and Control**, DAAC 2018
- CESARINI, D. Bartolini, A. and Benini L.. "Energy Saving and Thermal Management Opportunities in a Workload-Aware MPI Runtime for a Scientific HPC Computing Node." Parallel Computing is Everywhere 32 (2018): 277.
- Bartolini, A., Diversi, R., Cesarini, D., & Beneventi, F.. **Self-Aware Thermal Management for High Performance Computing Processors**. IEEE Design & Test.
- Borghesi, A., Bartolini, A., Lombardi, M., Milano, M., & Benini, L. (2018). **Anomaly Detection using Autoencoders in High Performance Computing Systems**. arXiv preprint arXiv:1811.05269. (to appear in IAAI19)